New Codes and Inner Bounds for Exact Repair in Distributed Storage Systems

Sreechakra Goparaju Salim El Rouayheb Robert Calderbank

arXiv:1402.2343v1 [cs.IT] 11 Feb 2014

Abstract

We study the exact-repair tradeoff between storage and repair bandwidth in distributed storage systems (DSS). We give new inner bounds for the tradeoff region and provide code constructions that achieve these bounds.

I. Introduction

The study of erasure codes which repair single node failures in a distributed storage system more efficiently was initiated by Dimakis et al in [1]. The codes, called *regenerating codes*, are constructed on a system of n storage nodes and store a file of size $\mathcal{M}$, such that the data stored in any k nodes is sufficient to recover the $\mathcal{M}$ symbols. Furthermore, each node stores α symbols and each failed node requires a transmission of γ symbols, termed as the *repair bandwidth*, by d other nodes, called *helper* nodes (d is also known as the *repair degree*). An optimal tradeoff between storage and repair-bandwidth (α and γ) is characterized in [1] for the so-called case of *functional repair*, where the repaired (or replaced) node need not store the exact copy of the data present in the failed node. This tradeoff region is given by the following expression:

$$\mathcal{M} \leq \sum_{i=0}^{k-1} \min\left\{\alpha, (d-i)\frac{\gamma}{d}\right\}. \tag{1}$$

However, when an exact duplication of data onto the repaired node, or *exact repair*, is required, this characterization remains an open problem in general. Since exact repair is a stricter constraint than functional repair, the tradeoff curve in (1) forms an outer bound to the exact-repair region for α and γ (for a given $\mathcal{M}, n, k$ and d). The extremal points of this tradeoff have attracted the most attention, being the optimal points in terms of α and γ. These are respectively called the *minimum storage regenerating* (MSR) point, with

$$\begin{aligned} \alpha &= \alpha_{\mathsf{msr}} = \frac{\mathcal{M}}{k}, \text{ and} \\ \gamma &= \gamma_{\mathsf{msr}} = \frac{d\mathcal{M}}{k(d-k+1)}, \end{aligned} \tag{2}$$

and the *minimum bandwidth regenerating* (MBR) point, with

$$\begin{aligned} \alpha &= \alpha_{\mathsf{mbr}} = \frac{2d\mathcal{M}}{k(2d-k+1)}, \text{ and} \\ \gamma &= \gamma_{\mathsf{mbr}} = \frac{2d\mathcal{M}}{k(2d-k+1)}, \end{aligned} \tag{3}$$

Both points have been shown to be achieveable for exact repair for all (n, k, d), using finite-length or asymptotic codes; see, for example, [2], [3], [4], and [5].

Recently, Tian [6] fully determined the optimal tradeoff for the case of an $(n, k, d) = (4, 3, 3)$ DSS. The optimal tradeoff regions for all other cases of (n, k, d) are still under investigation. Indeed, for some time, no inner bounds existed that were tighter than the space-sharing (or time-sharing) bound between the MSR and the MBR points. However, recently, new codes have been independently discovered, which improve upon the space-sharing bound, and lie in the so-called *intermediate* region, that is, between the MSR and the MBR points. In [7], block-designs

S. Goparaju is with the EE Department, Princeton University, USA (e-mail: goparaju@princeton.edu)
S. El Rouayheb is with the ECE Department, Illinois Institute of Technology, USA (e-mail: salim@iit.edu).
R. Calderbank is with the EE Department of Electrical Engineering, Duke University, USA (e-mail: robert.calderbank@duke.edu).

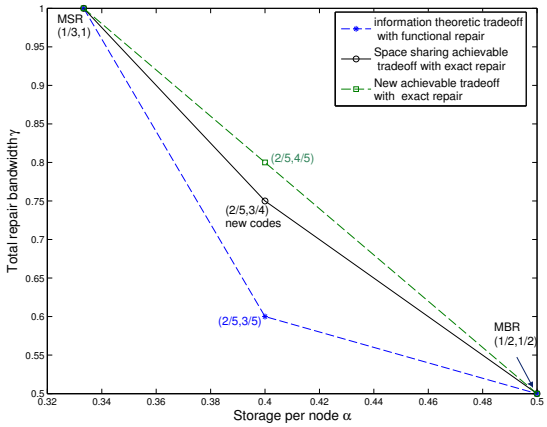

Fig. 1. Example of the new achievable exact-repair tradeoff for the $(5, 3, 3)$ case for a file size $\mathcal{M} = 1$. The code construction achieving the point $(\frac{2}{5}, \frac{3}{4})$ is described in Sec. II.

and erasure codes were used in a layered structure to build intermediate codes which are simple to implement. In [8], layering is used to construct codes first for the parameters (n, k, k), i.e., when $d = k$, and then extended to (n, k, d), when $d > k$. In [9], new (n, k, d) regenerating codes at intermediate points are formed using MSR points for lower parameters.

Contribution: We construct new intermediate codes which in turn lead to newer inner bounds for the storage/repair bandwidth tradeoff region for exact repair. As an example, for a $(5, 3, 3)$ DSS, our codes achieve the new point $(\alpha, \gamma) = (2/5, 3/4)$, which leads to the new achievable tradeoff shown in Fig. 1. This example is explained in detail in Section II, and is used to provide an intuition behind our code construction. Our technique is inspired by the constructions in [8] and [9], and the study of heterogeneous DSS in [10]. We use MSR code constructions for smaller codes to achieve intermediate points for larger parameters. We then generalize this technique in order to improve the repair bandwidth γ while keeping the rest of the parameters – n, k, d, α, and $\mathcal{M}$ – constant. In Section III, we describe our code construction and the ensuing tradeoff inner bound for the case when $d = k$. In Section IV, we extend this construction to the case when $d > k$. We then describe a second construction for this regime which uses ideas in [8] in a flavor similar to the way we improve the repair bandwidth. We show that this construction proves as useful as the first construction by plotting the tradeoff region for an example DSS; see Fig. 4. Finally, we conclude in Section V.

II. Primer: A $(5, 3, 3)$ Example

To illustrate our results, we start with a $(5, 3, 3)$ DSS consisting of five nodes $v_1, \ldots, v_5$. We give a new regenerating code for this DSS, which leads to an improved inner (achievable) bound on the storage/repair bandwidth tradeoff curve for this DSS under exact repair. The general results and proofs will be detailed in the subsequent sections.

The idea is to start with an MSR code for a smaller $(4, 2, 3)$ DSS with $(\hat{n}, \hat{k}) = (4, 2)$ consisting of the nodes $\{v_1, \ldots, v_4\}$ as in Fig. 2(a). Here, we pick the $(4, 2, 3)$ MSR code from [11]. Each node stores half a symbol $(\hat{\alpha} = 1/2)$, making the total stored file size $\mathcal{M} = 1$. We extend this code by adding an empty node v_5. The resulting DSS is heterogeneous since different nodes store different amounts of data. Moreover, the total repair bandwidth depends on which helper nodes participate in the repair process. There are two cases here.

Case 1: All the $d = 3$ helper nodes belong to the small MSR code, i.e., are chosen from $\{v_1, \ldots, v_4\}$. Here, we can achieve the optimal repair bandwidth given by the MSR code, i.e., the repair bandwidth is $3/4$.

Case 2: Only two nodes of the small MSR code help in the repair, the third node being v_5. In this case, the repair proceeds by downloading the whole file from the small code (and nothing from node v_5). The repair bandwidth is 1.

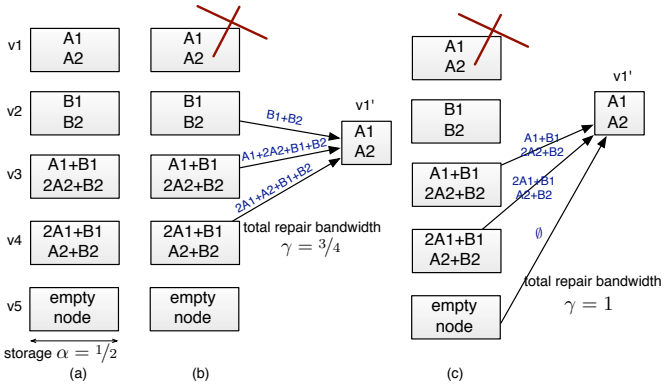

Fig. 2. Part (a) A $(5, 3, 3)$ code constructed by appending a $(4, 2, 3)$ MSR code [11] to an empty node. Part (b) shows the repair of a non-empty node in the $(5, 3, 3)$ system when the $d = 3$ helper nodes belong to the MSR code. Part (c) shows the repair when the d helper nodes consist of only 2 nodes in the MSR code.

From [10], we know that we can construct a homogeneous code by "glueing" together all the $n! = 120$ permuted copies of the heterogeneous code. This glued code achieves the average storage per node α and average repair bandwidth γ of the heterogeneous code. For the code in Fig. 2, a fraction of $4/5$ of the nodes are non-empty. Furthermore, for the repair of a non-empty node, a fraction of $1/4$ of helper sets consists of three helper nodes (*Case 1*) in the $(4, 2, 3)$ code. Thus, we have

$$\alpha \;=\; \frac{2}{5}, \quad \text{and} \quad \gamma_i \;=\; \frac{15}{16}, \;\text{for}\; i \in \{1, 2, 3, 4\},$$

where γ_i is the average repair bandwidth when repairing node v_i, $i = 1, \ldots, 5$. The repair bandwidth γ_5 when node v_5 fails is zero since no data needs to be downloaded. Therefore, the average total repair bandwidth, over all failures and all choices of helper nodes, is $\gamma = 3/4$.

The tradeoff curve for functional repair [1] is given by the piecewise linear curve joining the MSR point $(\alpha, \gamma) = (1/3, 1)$, an intermediate point $(2/5, 3/5)$ and the MBR point $(1/2, 1/2)$ (see Fig. 1). The MSR and MBR points are also achievable with exact repair [4], [2]. By space sharing, we can achieve any point on the straight line joining the MSR and MBR points. The code constructions in [9] and [8] both correspond to the point $(2/5, 4/5)$ on the space sharing curve for the $(5, 3, 3)$ example. The proposed code lies beneath this curve in Fig. 1 and thus leads to an improved achievable tradeoff curve for exact repair.

III. Code Construction for $d = k$

Our goal is to build (n, k, d) regenerating codes with exact repair for intermediate points on the tradeoff curve. We focus first on the case when $d = k$ and address the case when $d > k$ in the next section. We follow the method in [9] and build the (n, k, d) code from a small $(\hat{n}, \hat{k})$ code to which we append $n - \hat{n}$ empty nodes as depicted in Fig 3. In the previous example, we had $(n, k, d) = (5, 3, 3)$ and $(\hat{n}, \hat{k}) = (4, 2)$. Notice that we do not specify the repair degree $\hat{d}$ yet.

A. Construction

The key ingredient in our construction is to pick the small code to be an MSR code with minimum repair bandwidth for all possible values of the repair degree $\hat{d} = \hat{k}, \hat{k} + 1, \ldots, \hat{n} - 1$, simultaneously, that is, the repair bandwidth $\hat{\gamma}_{\hat{d}}$ when the repair degree is $\hat{d}$ is given by

$$\hat{\gamma}_{\hat{d}} \;=\; \frac{\hat{d}}{\hat{d} - \hat{k} + 1}\,\hat{\alpha}, \tag{4}$$

where $\hat{\alpha}$ is the storage per node in the small code; see (2). Constructions of MSR codes that can asymptotically achieve the minimum repair bandwidth for any repair degree were described in [4] and are based on the interference

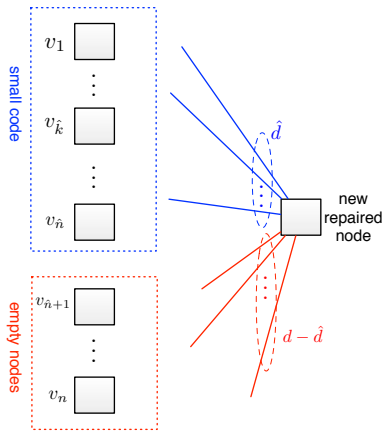

Fig. 3. General code construction.

alignment method. This code property, referred to as *universality* in [4], has been studied in depth in [12], where it is referred to as *opportunistic repair*. Note that for the special case of $\hat{n}-1=\hat{k}+1$, as in the example, any MSR code has this property since for $\hat{d}=\hat{n}-1=\hat{k}+1$, optimal repair bandwidth follows from the definition of an MSR code, and for $\hat{d}=\hat{k}$, the optimal repair requires downloading the whole file.

Hence we do not need to worry about $\hat{d}$ when choosing the parameters of the small code (as long as $\hat{n}-1\geq\hat{k}$, which is true for a non-trivial small code). Since file reconstruction and exact repair in the big code are inherited from the small code, we need to pick $\hat{k}$ to make sure that among the k nodes from $\{v_1,\ldots,v_n\}$ contacted by a user, there are always $\hat{k}$ nodes from $\{v_1,\ldots,v_{\hat{n}}\}$. This is achieved by choosing $\hat{n}$ and $\hat{k}$ such that the big code and small code have the same number of parity nodes, i.e.,

$$n-k=\hat{n}-\hat{k}. \tag{5}$$

Also, since $d=k$, the repair process is guaranteed to contact at least $\hat{k}$ nodes from $\{v_1,\ldots,v_{\hat{n}}\}$, and thus can be inherited from the small code. We vary $\hat{k}$ to take all the integer values between 1 and k while choosing $\hat{n}=\hat{k}+(n-k)$. Each value of $\hat{k}$ then results in a distinct regenerating code.

Ultimately, the code is formed of all the $n!$ permuted copies of the code in Fig. 3 glued together[1] as in [10]. This code achieves the point (α,γ), where α and γ are the average node storage capacity and repair bandwidth of the code in Fig. 3, respectively. Since the storage per node in the small code is $\hat{\alpha}$, we obtain

$$\alpha \;=\; \frac{\hat{n}}{n}\,\hat{\alpha}. \tag{6}$$

Let γ_1 be the average repair bandwidth when repairing a non-empty node. Among all the possible sets of helper nodes, there is a fraction $P_{\hat{d}}$ (defined in (7)) of sets that consist of exactly $\hat{d}$ nodes in the small code, where $\hat{d}$ ranges from $\hat{k}$ to $\hat{d}_{\mathsf{max}}$ and $\hat{d}_{\mathsf{max}}=\min\{\hat{n}-1,d\}$. So we obtain

$$\begin{aligned}\gamma_1 \;&=\; \sum_{\hat{d}=\hat{k}}^{\hat{d}_{\mathsf{max}}}\hat{\gamma}_{\hat{d}}P_{\hat{d}},\\ &=\; \sum_{\hat{d}=\hat{k}}^{\hat{d}_{\mathsf{max}}}\left(\frac{\hat{d}}{\hat{d}-\hat{k}+1}\right)\frac{\binom{\hat{n}-1}{\hat{d}}\binom{n-\hat{n}}{d-\hat{d}}}{\binom{n-1}{d}}\,\hat{\alpha}.\end{aligned} \tag{7}$$

[1]It may be possible to make the system homogeneous without using all the $n!$ permutations, but we don't worry about this issue here.

Because the average repair bandwidth when repairing an empty node is zero, the overall average repair bandwidth γ is given by

$$\gamma \;=\; \sum_{\hat{d}=\hat{k}}^{\hat{d}_{\mathsf{max}}} \left(\frac{\hat{d}}{\hat{d}-\hat{k}+1}\right) \frac{\binom{\hat{n}-1}{\hat{d}}\binom{n-\hat{n}}{d-\hat{d}}}{\binom{n-1}{d}}\,\alpha. \tag{8}$$

Since the small code is MSR, the stored file can be of size up to

$$\mathcal{M} \;=\; \hat{k}\hat{\alpha}. \tag{9}$$

B. Inner Bound

From the above construction and from (6), (8), and (9), we have the following new achievable exact-repair tradeoff. Since it is possible that the points obtained by the construction do not enclose the MBR and the MSR points, we also add them to the tradeoff region.

Theorem 1: There exist (n,k,d) exact-repair regenerating codes, with $d=k$, that store a file of size $\mathcal{M}$ and achieve any point in the region $\mathsf{con}(\mathcal{R})$, the convex hull of $\mathcal{R}$, where

$$\mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_{\mathsf{mbr}} \cup \mathcal{R}_{\mathsf{msr}}, \tag{10}$$

$$\mathcal{R}_1 = \bigcup_{\hat{k}=1}^{k} \left\{(\bar{\alpha},\bar{\gamma}) : \bar{\alpha} \geq \alpha(\hat{k}), \bar{\gamma} \geq \gamma(\hat{k})\right\}, \tag{11}$$

$$\mathcal{R}_{\mathsf{mbr}} = \left\{(\bar{\alpha},\bar{\gamma}) : \bar{\alpha} \geq \alpha_{\mathsf{mbr}}, \bar{\gamma} \geq \gamma_{\mathsf{mbr}}\right\}, \quad \text{and} \tag{12}$$

$$\mathcal{R}_{\mathsf{msr}} = \left\{(\bar{\alpha},\bar{\gamma}) : \bar{\alpha} \geq \alpha_{\mathsf{msr}}, \bar{\gamma} \geq \gamma_{\mathsf{msr}}\right\}. \tag{13}$$

Here $\alpha(\hat{k})$ and $\gamma(\hat{k})$ are given by (6) and (8), respectively. Notice that $\hat{k}$ also determines $\hat{n}$ from (5), and $\hat{\alpha}$ is determined by (9).

Remark 1: For the case of $d=k$, the two code constructions in [9], [8] achieve the same points on the tradeoff, that are obtained when the summation in (8) contains only the term corresponding to $\hat{d}=\hat{k}$. The inner bound in Theorem 1 is thus an improvement whenever the summation in (8) has more than one term. This happens for all system parameters except when $k=d=n-1$, for which the three inner bounds coincide.

Remark 2: In general, we only know of asymptotic schemes [4] which achieve the repair bandwidth given in (4) for all feasible values of $\hat{d}$. However, for $n-k=2$, the summation in (8) has summands corresponding to the two extreme cases of $\hat{d}=\hat{k}$ and $\hat{d}=\hat{k}+1$, where the former requires each helper disk to transmit all its data as repair bandwidth, and the latter requires each helper disk to transmit only $1/(n-k)$ of its data. The latter is the optimal repair bandwidth for an $(\hat{n},\hat{k},\hat{n}-1)$ MSR code, for which several finite code constructions exist, e.g. [2], [3] and [5].

IV. Code Constructions for $d > k$

The construction in Section III can be generalized to the case when $d>k$ by using the same small code as before. The difference however arises during the repair of a failed node in the big code, as detailed in *Construction 1* below. A second construction, *Construction 2*, can be achieved using similar ideas, this time following the method in [8].

A. Construction 1

As before, we construct the big (n,k,d) regenerating code using the permuted copies of a small MSR $(\hat{n},\hat{k})$ code, where $\hat{n}-\hat{k}=n-k$, and which is repair-bandwidth-optimal for all possible degrees of repair $\hat{d} \in \{\hat{k},\hat{k}+1,\ldots,\hat{n}-1\}$.

Upon failure of a non-empty node, say v_1 (see Fig. 3), in the (n,k,d) regenerating code, a fraction $P_{\hat{d}}$ of helper sets consist of exactly $\hat{d}$ nodes in the small code. The repair bandwidth γ_1 is then given by

$$\gamma_1 = \sum_{\hat{d}=\hat{d}_{\mathsf{min}}}^{\hat{d}_{\mathsf{max}}} \hat{\gamma}_{\hat{d}} P_{\hat{d}}, \tag{14}$$

$$= \sum_{\hat{d}=\hat{d}_{\mathsf{min}}}^{\hat{d}_{\mathsf{max}}} \left(\frac{\hat{d}}{\hat{d}-\hat{k}+1}\right) \frac{\binom{\hat{n}-1}{\hat{d}}\binom{n-\hat{n}}{d-\hat{d}}}{\binom{n-1}{d}} \hat{\alpha}, \tag{15}$$

where $\hat{d}_{\mathsf{min}} = d-(n-\hat{n})$, and $\hat{d}_{\mathsf{max}} = \min\{\hat{n}-1, d\}$ as defined earlier. Notice that unlike when $d=k$, any d (helper) nodes must contain at least $d-(n-\hat{n})$ nodes in the small code, where it can be verified that $\hat{d}_{\mathsf{min}} > \hat{k}$.

Since the repair bandwidth γ_i for the failure of an empty node v_i is 0, the overall repair bandwidth γ is given by

$$\gamma = \sum_{\hat{d}=\hat{d}_{\mathsf{min}}}^{\hat{d}_{\mathsf{max}}} \left(\frac{\hat{d}}{\hat{d}-\hat{k}+1}\right) \frac{\binom{\hat{n}-1}{\hat{d}}\binom{n-\hat{n}}{d-\hat{d}}}{\binom{n-1}{d}} \alpha. \tag{16}$$

The stored file size, as before, is given by $\mathcal{M} = \hat{k}\hat{\alpha}$, and the average storage per node is given by $\alpha = (\hat{n}/n)\hat{\alpha}$.

B. Construction 2

When $d > k$, an alternative code can be constructed by viewing the big (n,k,d) regenerating code as an (n,d,d) regenerating code as obtained by *Construction 1* (in Section III), and calculating the *amount of information*[2] that any set of k nodes contains. This quantity $\mathcal{M}_k$ is the same for any set of k nodes from the symmetry in *Construction 1*. It can be further verified that $\mathcal{M}_k$ is equal to the average amount of information contained in all the possible sets of size k (in the heterogeneous code). Let $\mathcal{M}_d$ ($\geq \mathcal{M}_k$), be the file size of the (n,d,d) code obtained from the construction, and let $\mathcal{M}_k$ be the size of the file that we wish to store in the DSS. The underlying motivation is to enable the recovery of the file from any k nodes (which as we defined contain $\mathcal{M}_k$ amount of information). One way to achieve this is to concatenate a maximum rank distance (MRD) code and the (n,d,d) code. To do so, the MRD codeword is constructed using a linearized polynomial with coefficients as the $\mathcal{M}_k$ symbols of our file. The output codeword of this code can then be viewed as a set of vectors which are evaluations of this polynomial on $\mathcal{M}_d$ points in a specific field, which are then fed as the input message file for the (n,d,d) regenerating code constructed using *Construction 1*. The details of this concatenation are skipped here but can be found, for example, in [8]. We now calculate the overall file size $\mathcal{M}$ ($=\mathcal{M}_k$) achievable.

Suppose that the small code being used is an $(\hat{n},\hat{k})$ MSR code, where now $\hat{n}-\hat{k} = n-d$. Let this code have $\hat{\alpha}$ units of storage per node and be repair-bandwidth optimal for all possible degrees of repair $\hat{d} \in \{\hat{k},\hat{k}+1,\ldots,\hat{n}-1\}$. Consider any set $\mathcal{K}$ of k nodes in the big code. We noticed in Section III when $d=k$ that $\mathcal{K}$ contains at least $\hat{k}$ nodes of the small code in all permutations of the small code and $(n-\hat{n})$ empty nodes. When $d>k$, we use the property that the amount of information (entropy or linear dimension for the case of linear codes) contained in a set of ω nodes in the $(\hat{n},\hat{k})$ MSR code is given by $H_\omega = \min\{\omega,\hat{k}\}\,\hat{\alpha}$.

Therefore, if Q_ω is the fraction of permutations for which ω (and only ω) of the nodes in the small code occur

[2]This can be viewed in different forms – rank of the code generating matrix punctured at the columns that correspond to the $(n-k)$ nodes that are not selected, as in [8], or as the entropy of the random variables corresponding to the k selected nodes.

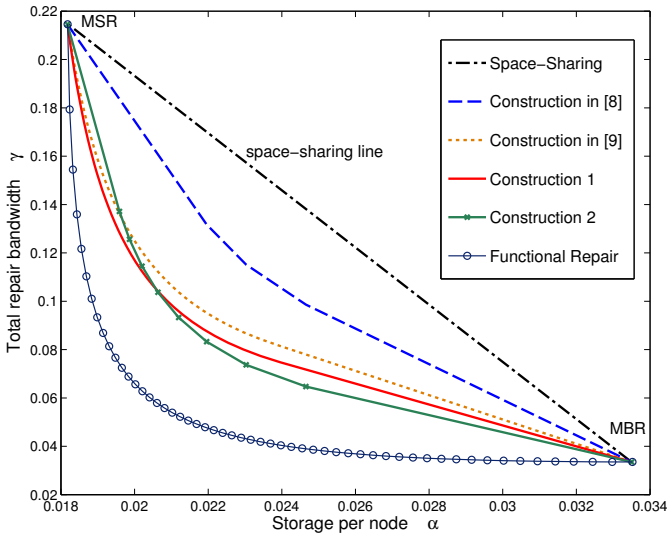

Fig. 4. Comparison among the different bounds existing in the literature and the new inner bounds achieved by constructions 1 and 2 in Section IV for the case of $(n, k, d) = (61, 55, 59)$. The overall tradeoff region as defined in Theorem 2 is achieved by the convex hull of the two new constructions.

in $\mathcal{K}$, the average amount of information $\mathcal{M}_k$ in $\mathcal{K}$ is given by

$$\mathcal{M}_k = \sum_{\omega=\omega_{\min}}^{\omega_{\max}} H_\omega Q_\omega, \tag{17}$$

$$= \sum_{\omega=\omega_{\min}}^{\omega_{\max}} \min\{\omega, \hat{k}\} \frac{\binom{\hat{n}}{\omega}\binom{n-\hat{n}}{k-\omega}}{\binom{n}{k}} \hat{\alpha}, \tag{18}$$

where $\omega_{\min} = \max\{1, \hat{n} - (n-k)\}$ and $\omega_{\max} = \max\{k, \hat{n}\}$ give the limits determined by the possible overlap sizes of the small code and $\mathcal{K}$.

Observe that the repair process is not disrupted during this process and hence the average repair bandwidth remains the same as in (8).

C. Inner Bound

A natural question arises whether one of the constructions for $d > k$ is better than the other. We show using an example that this is not necessarily true.

Example: Consider[3] an $(n, k, d) = (61, 55, 59)$ DSS. Fig. 4 plots the inner bounds achieved for the storage/repair-bandwidth tradeoff using different code constructions. We see that for this example, the inner bound achieved by using the codes in [9] strictly encompasses the bound achieved by those in [8]. However, *Construction 1* and *Construction 2*, do not exhibit such a relationship. Whereas *Construction 1* outperforms *Construction 2* at points closer to the MSR point, *Construction 2* generates a tighter inner bound near the MBR point.

It follows therefore that an overall inner bound must incorporate points from both constructions, as formalized in the next theorem.

[3]We considered this example for easy comparison with the adjacent case of $(61, 55, 60)$ considered in [8].

Theorem 2: There exist (n, k, d) exact-repair regenerating codes that store a file of size $\mathcal{M} = 1$ and achieve any point in the region $\mathsf{con}(\mathcal{R})$, the convex hull of $\mathcal{R}$, where

$$\mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_{\mathsf{mbr}} \cup \mathcal{R}_{\mathsf{msr}}, \tag{19}$$

$$\mathcal{R}_1 = \bigcup_{\hat{k}_1=1}^{k} \left\{ (\bar{\alpha}, \bar{\gamma}) : \bar{\alpha} \geq \frac{\alpha(\hat{k}_1)}{\mathcal{M}(\hat{k}_1)}, \bar{\gamma} \geq \frac{\gamma(\hat{k}_1)}{\mathcal{M}(\hat{k}_1)} \right\}, \tag{20}$$

$$\mathcal{R}_2 = \bigcup_{\hat{k}_2=1}^{d} \left\{ (\bar{\alpha}, \bar{\gamma}) : \bar{\alpha} \geq \frac{\alpha(\hat{k}_2)}{\mathcal{M}_k(\hat{k}_2)}, \bar{\gamma} \geq \frac{\gamma(\hat{k}_2)}{\mathcal{M}_k(\hat{k}_2)} \right\}, \tag{21}$$

$$\mathcal{R}_{\mathsf{mbr}} = \left\{ (\bar{\alpha}, \bar{\gamma}) : \bar{\alpha} \geq \frac{\alpha_{\mathsf{mbr}}}{\mathcal{M}}, \bar{\gamma} \geq \frac{\gamma_{\mathsf{mbr}}}{\mathcal{M}} \right\}, \quad \text{and} \tag{22}$$

$$\mathcal{R}_{\mathsf{msr}} = \left\{ (\bar{\alpha}, \bar{\gamma}) : \bar{\alpha} \geq \frac{\alpha_{\mathsf{msr}}}{\mathcal{M}}, \bar{\gamma} \geq \frac{\gamma_{\mathsf{msr}}}{\mathcal{M}} \right\}, \tag{23}$$

where $\mathcal{R}_1$ is the tradeoff region corresponding to *Construction* 1, and $\alpha(\hat{k}_1), \gamma(\hat{k}_1)$ and $\mathcal{M}(\hat{k}_1)$ in (20) are given by (6), (16), and (9), respectively, evaluated at $\hat{k} = \hat{k}_1$ and $\hat{n} = n + \hat{k}_1 - k$. Similarly, $\mathcal{R}_2$ corresponds to *Construction 2*, and $\alpha(\hat{k}_2), \gamma(\hat{k}_2)$ and $\mathcal{M}_k(\hat{k}_2)$ in (21) are given by (6), (8) and (18), respectively, evaluated at $\hat{k} = \hat{k}_2$ and $\hat{n} = n + \hat{k}_1 - d$. The regions corresponding to the MBR and the MSR points are evaluated at the given k and d.

V. Conclusion

Determining the achievable region in the tradeoff between storage and exact repair bandwidth for distributed storage systems is an important problem that is still open in general, with this region fully characterized only for the $(4, 3, 3)$ case [6]. This paper makes a contribution towards solving this problem by proposing new constructions of regenerating codes with exact repair that achieve new points in the tradeoff region leading to improved inner bounds. Our code constructions are a generalization of the codes in [9] and [8]. The main idea is to construct new codes using already known code constructions for systems with smaller parameters. Our constructions have two key ingredients: (i) Use as a building block a minimum storage code (MSR) with optimal repair bandwidth for all possible repair degrees, and (ii) Allow a heterogeneous repair bandwidth that depends on the nodes contacted for repair. The code is then "homogenized" by appending together all of its $n!$ permutations as described in [10]. This method leads to new achievable points on the tradeoff and result in tighter inner bounds for all systems having more than one parity node ($n > k + 1$).

References

[1] A. G. Dimakis, P. G. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network Coding for Distributed Storage Systems," in *IEEE Transactions on Information Theory*, vol. 56, Sep. 2010, pp. 4539–4551.

[2] K. V. Rashmi, N. B. Shah, and P. V. Kumar, "Optimal Exact-Regenerating Codes for Distributed Storage at the MSR and MBR Points via a Product-Matrix Construction," in *IEEE Transactions on Information Theory*, vol. 57, Aug. 2011, pp. 5227–5239.

[3] D. Papailiopoulos, A. Dimakis, and V. Cadambe, "Repair Optimal Erasure Codes through Hadamard Designs," in *Proceedings of the 49th Annual Allerton Conference on Communication, Control, and Computing*, Sep. 2011, pp. 1382–1389.

[4] V. Cadambe, S. Jafar, H. Maleki, K. Ramchandran, and C. Suh, "Asymptotic Interference Alignment for Optimal Repair of MDS codes in Distributed Storage," in *IEEE Transactions on Information Theory*, vol. 59, May 2013, pp. 2974–2987.

[5] Z. Wang, I. Tamo, and J. Bruck, "On Codes for Optimal Rebuilding Access," in *Proceedings of the 49th Annual Allerton Conference on Communication, Control, and Computing*, 2011, pp. 1374–1381.

[6] C. Tian, "Rate Region of the $(4, 3, 3)$ Exact-Repair Regenerating Codes," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, Jul. 2013, pp. 1426–1430.

[7] C. Tian, V. Aggarwal, and V. A. Vaishampayan, "Exact-Repair Regenerating Codes via Layered Erasure Correction and Block Designs," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, 2013, pp. 1431–1435.

[8] B. Sasidharan and P. V. Kumar, "High-Rate Regenerating Codes Through Layering," in *arxiv.org*, Mar. 2013. [Online]. Available: http://arxiv.org/abs/1301.6157

[9] T. Ernvall, "Exact-regenerating codes between mbr and msr points," in *arXiv:1304.5357v1*, 2013.

[10] T. Ernvall, S. El Rouayheb, C. Hollanti, and H. V. Poor, "Capacity and Security of Heterogeneous Distributed Storage Systems," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, 2013, pp. 1247–1251.

[11] C. Suh and K. Ramchandran, "Exact-Repair MDS Codes for Distributed Storage using Interference Alignment," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, Jun. 2010, pp. 161–165.

[12] V. Aggarwal, C. Tian, V. A. Vaishampayan, and Y. R. Chen, "Distributed Data Storage Systems with Opportunistic Repair," in *arxiv.org*, Nov. 2013. [Online]. Available: http://arxiv.org/abs/1311.4096